

ABSTRACT OF THE DISCLOSURE

The present invention provides a technique for synchronisation between pipelines in a data processing apparatus. The data processing apparatus comprises a main processor operable to execute a sequence of instructions, the main processor comprising a first pipeline having a first plurality of pipeline stages, and a coprocessor operable to execute coprocessor instructions in said sequence of instructions. The coprocessor comprises a second pipeline having a second plurality of pipeline stages, and each coprocessor instruction is arranged to be routed through both the first pipeline and the second pipeline. Furthermore, at least one synchronising queue is provided coupling a predetermined pipeline stage in one of the pipelines with a partner pipeline stage in the other of the pipelines, the predetermined pipeline stage being operable to cause a token to be placed in the synchronising queue when processing a coprocessor instruction, and the partner pipeline stage being operable to process that coprocessor instruction upon receipt of the token from the synchronising queue. By this approach, the first and second pipelines are synchronised between the predetermined pipeline stage and the partner pipeline stage, and hence ensures that the pipelines are correctly synchronised for crucial transfers of information without requiring that strict synchronisation at all stages is necessary.